arXiv:1404.6999v1 [cs.AI] 28 Apr 2014

Preliminary Report on WASP 2.0*

Mario Alviano, Carmine Dodaro and Francesco Ricca
Department of Mathematics and Computer Science, University of Calabria, Italy
{alviano,dodaro,ricca}@mat.unical.it

Abstract

Answer Set Programming (ASP) is a declarative programming paradigm. The intrinsic complexity of the evaluation of ASP programs makes the development of more effective and faster systems a challenging research topic. This paper reports on the recent improvements of the ASP solver WASP. WASP is undergoing a refactoring process which will end up in the release of a new and more performant version of the software. In particular the paper focus on the improvements to the core evaluation algorithms working on normal programs. A preliminary experiment on benchmarks from the 3rd ASP competition belonging to the NP class is reported. The previous version of WASP was often not competitive with alternative solutions on this class. The new version of WASP shows a substantial increase in performance.

Introduction

Answer Set Programming (ASP) (Gelfond and Lifschitz 1991) is a declarative programming paradigm which has been proposed in the area of non-monotonic reasoning and logic programming. The idea of ASP is to represent a given computational problem by a logic program whose answer sets correspond to solutions, and then use a solver to find them.

Despite the intrinsic complexity of the evaluation of ASP, after twenty years of research many efficient ASP systems have been developed. (e.g. (Alviano et al. 2011; Gebser et al. 2007; Lierler and Maratea 2004)). The availability of robust implementations made ASP a powerful tool for developing advanced applications in the areas of Artificial Intelligence, Information Integration, and Knowledge Management. These applications of ASP have confirmed the viability of the use of ASP. Nonetheless, the interest in developing more effective and faster systems is still a crucial and challenging research topic, as witnessed by the results of the ASP Competition series (see e.g. (Calimeri, Ianni, and Ricca 2014)).

*This research has been partly supported by the European Commission, European Social Fund of Regione Calabria, the Regione Calabria under project PIA KnowRex POR FESR 2007- 2013 BURC n. 49 s.s. n. 1 16/12/2010, and the Italian Ministry of University and Research under PON project "Ba2Know S.I.-LAB" n. PON03PE_0001.

This paper reports on the recent improvements of the ASP solver for propositional programs WASP (Alviano et al. 2013). The new version of WASP is inspired by several techniques that were originally introduced for SAT solving, like the Davis-Putnam-Logemann-Loveland (DPLL) backtracking search algorithm (Davis, Logemann, and Loveland 1962), *clause learning* (Zhang et al. 2001), *backjumping* (Gaschnig 1979), *restarts* (Gomes, Selman, and Kautz 1998), and *conflict-driven heuristics* (Moskewicz et al. 2001). The mentioned SAT-solving methods have been adapted and combined with state-of-the-art pruning techniques adopted by modern native ASP solvers (Alviano et al. 2011; Gebser et al. 2007). In particular, the role of Boolean Constraint Propagation in SAT-solvers is taken by a procedure combining the *unit propagation* inference rule with inference techniques based on ASP program properties. In particular, support inferences are implemented via Clark's completion, and the implementation of the well-founded operator is based on source pointers (Simons, Niemelä, and Soininen 2002).

In the following, we overview the techniques implemented by the 2.0 version of WASP, focusing on the improvements to the core evaluation algorithms working on normal programs. Then we compare the new implementation with the previous one.

We also report on a preliminary experiment in which we compare the old and new versions of WASP with the latest version of clasp, which is the solver that won the 3rd and 4th edition of the ASP competition. Benchmarks were taken from the 3rd ASP competition and belong to the NP class, i.e., the class of problems where the previous version of WASP was often not competitive with alternative solutions. The result show that WASP 2.0 is substantially faster than WASP 1.0 and is often competitive with clasp.

ASP Language

Let $\mathcal{A}$ be a countable set of propositional atoms. A *literal* is either an atom (a positive literal), or an atom preceded by the *negation as failure* symbol $\sim$ (a negative literal). The complement of a literal ℓ is denoted $\overline{\ell}$, i.e., $\overline{a} = \sim a$ and $\overline{\sim a} = a$ for an atom a. This notation extends to sets of literals, i.e., $\overline{L} := \{\overline{\ell} \mid \ell \in L\}$ for a set of literals L.

A *program* is a finite set of rules of the following form:

$$a_0 \texttt{ :- } a_1, \ldots, a_m, \sim a_{m+1}, \ldots, \sim a_n \qquad (1)$$

where $n \geq m \geq 0$ and each a_i ($i = 0, \ldots, n$) is an atom. The atom a_0 is called head, and the conjunction $a_1, \ldots, a_m, \sim a_{m+1}, \ldots, \sim a_n$ is referred to as body. Rule r is said to be regular if $H(r) \neq \bot$, where $\bot$ is a fixed atom in $\mathcal{A}$, and a constraint otherwise. For a rule r of the form (1), the following notation is also used: $H(r)$ denotes the head atom a_0; $B(r)$ denotes the set $\{a_1, \ldots, a_m, \sim a_{m+1}, \ldots, \sim a_n\}$ of body literals; $B^+(r)$ and $B^-(r)$ denote the set of atoms appearing in positive and negative body literals, respectively; $C(r) := H(r) \cup \overline{B(r)}$ is the clause representation of r.

An *interpretation* I is a set of literals, i.e., $I \subseteq \mathcal{A} \cup \overline{\mathcal{A}}$. Intuitively, literals in I are true, literals whose complements are in I are false, and all other literals are undefined. I is total if there are no undefined literals, and I is inconsistent if $\bot \in I$ or there is $a \in \mathcal{A}$ such that $\{a, \sim a\} \subseteq I$. An interpretation I satisfies a rule r if $C(r) \cap I \neq \emptyset$, while I violates r if $C(r) \subseteq \overline{I}$. A *model* of a program $\mathcal{P}$ is a consistent, total interpretation satisfying all rules of $\mathcal{P}$. The semantics of a program $\mathcal{P}$ is given by the set of its *answer sets* (or stable models) (Gelfond and Lifschitz 1991), where an interpretation I is an answer set for $\mathcal{P}$ if I is a subset-minimal model of the reduct $\mathcal{P}^I$ obtained by deleting from $\mathcal{P}$ each rule r such that $B^-(r) \cap I \neq \emptyset$, and then by removing all the negative literals from the remaining rules.

Answer Set Computation in WASP 2.0

In this section we review the algorithms implemented in WASP 2.0. The presentation is properly simplified to focus on the main principles.

Completion and Program Simplification

The first step of the evaluation in WASP 2.0 is a program transformation step. The input program first undergoes a Clark's completion transformation step, and then is simplified applying techniques in the style of satelite (Eén and Biere 2005). Given a rule $r \in \mathcal{P}$, let aux_r denote a fresh atom, i.e., an atom not appearing elsewhere. The completion of $\mathcal{P}$, denoted $Comp(\mathcal{P})$, consists of the following clauses:

- $\{\sim a, aux_{r_1}, \ldots, aux_{r_n}\}$ for each atom a occurring in $\mathcal{P}$, where $r_1, \ldots, r_n$ are the rules of $\mathcal{P}$ whose head is a;
- $\{H(r), \sim aux_r\}$ and $\{aux_r\} \cup \overline{B(r)}$ for each rule $r \in \mathcal{P}$;
- $\{\sim aux_r, \ell\}$ for each $r \in \mathcal{P}$ and $\ell \in B(r)$.

After the computation of Clark's completion, *simplification* techniques are applied (Eén and Biere 2005). These consist of polynomial algorithms for strengthening and for removing redundant clauses, and also include atoms elimination by means of clause rewriting.

Main Algorithm

An answer set of a given propositional program $Comp(\mathcal{P})$ is computed in WASP 2.0 by using Algorithm 1, which is similar to the DPLL procedure in SAT solvers. Initially, interpretation I is set to $\{\sim\bot\}$. Function Propagate (line 2) extends

Algorithm 1: Compute Answer Set

```
   Input  : An interpretation I for a program Comp(P)
   Output: An answer set for Comp(P) or Incoherent
1  begin
2      while Propagate(I) do
3          if I is total then
4              return I;
5          ℓ := ChooseUndefinedLiteral();
6          I' := ComputeAnswerSet(I ∪ {ℓ});
7          if I' ≠ Incoherent then
8              return I';
9          if there are violated (learned) clauses then
10             return Incoherent;
11     AnalyzeConflictAndLearnClauses(I);
12     return Incoherent;
```

Function Propagate(I)

```
1  while UnitPropagation(I) do
2      if not WellFoundedPropagation(I) then
3          return true;
4  return false;
```

I with those literals that can be deterministically inferred. This function returns false if an inconsistency (or conflict) is detected, true otherwise. When no inconsistency is detected, interpretation I is returned if total (lines 2–3). Otherwise, an undefined literal, say ℓ, is chosen according to some heuristic criterion (line 5). Then computation then proceeds with a recursive call to ComputeAnswerSet on $I \cup \{\ell\}$ (line 6). In case the recursive call returns an answer set, the computation ends returning it (lines 7–8). Otherwise, the algorithm unrolls choices until consistency of I is restored (backjumping; lines 9–10), and the computation resumes by propagating the consequences of the clause learned by the conflict analysis. Conflicts detected during propagation are analyzed by procedure *AnalyzeConflictAndLearnClauses* (line 11).

The main algorithm is usually complemented with some heuristic techniques that control the number of learned clauses (which may be exponential in number), and possibly restart the computation to explore different branches of the search tree. Moreover, a crucial role is played by the heuristic criteria used for selecting branching literals. WASP 2.0 adopts the same branching and deletion heuristics of the SAT solver MiniSAT (Eén and Sörensson 2003). The restart policy is based on the sequence of thresholds introduced in (Luby, Sinclair, and Zuckerman 1993).

Propagation and clause learning are described in more detail in the following.

Propagation. WASP 2.0 implements two deterministic inference rules for pruning the search space during answer set computation. These propagation rules are named *unit* and *well-founded*. Unit propagation is applied first (line 1

of function Propagate). It returns false if an inconsistency arises. Otherwise, well-founded propagation is applied (line 2). Function WellFoundedPropagation may learn an implicit clause in P, in which case true is returned and unit propagation is applied on the new clause. When no new clause can be learned by WellFoundedPropagation, function Propagate returns true to report that no inconsistency has been detected. More in details, unit propagation is as in SAT solvers: An undefined literal ℓ is inferred by unit propagation if there is a rule r that can be satisfied only by ℓ, i.e., r is such that $\ell \in C(r)$ and $C(r) \setminus \{\ell\} \subseteq \overline{I}$. Concerning well-founded propagation, we must first introduce the notion of unfounded set. A set X of atoms is *unfounded* if for each rule r such that $H(r) \cap X \neq \emptyset$, at least one of the following conditions is satisfied: (i) $\overline{B(r)} \cap I \neq \emptyset$; (ii) $B^+(r) \cap X \neq \emptyset$; (iii) $I \cap H(r) \setminus X \neq \emptyset$. Intuitively, atoms in X can have support only by themselves. When an unfounded set X is found, function WellFoundedPropagation learns a clause forcing falsity of an atom in X. Clauses for other atoms in X will be learned on subsequent calls to the function, unless an inconsistency arises during unit propagation. In case of inconsistencies, indeed, the unfounded set X is recomputed.

Conflict Analysis and Learning. Clause learning acquires information from conflicts in order to avoid exploring the same search branch several times. WASP 2.0 adopts a learning schema based on the concept of the first Unique Implication Point (UIP) (Moskewicz et al. 2001), which is computed by analyzing the so-called implication graph. Roughly, the implication graph contains a node for each literal in I, and arcs from ℓ_i to ℓ_0 ($i = 1, \ldots, n$; $n \geq 1$) if literal ℓ_0 is inferred by unit propagation on clause $\{\ell_0, \ldots, \ell_n\}$. Each literal $\ell \in I$ is associated with a *decision level*, corresponding to the depth nesting level of the recursive call to ComputeAnswerSet on which ℓ is added to I. A node n in the implication graph is a UIP for a decision level d if all paths from the choice of level d to the conflict literals pass through n. The first UIP is the UIP for the decision level of the conflict that is closest to the conflict. The learning schema is as follows: Let u be the first UIP. Let L be the set of literals different form u occurring in a path from u to the conflict literals. The learned clause comprises u and each literal ℓ such that the decision level of ℓ is lower than the one of u and there is an arc $(\overline{\ell}, \ell')$ in the implication graph for some $\ell' \in L$.

Comparing WASP 1.0 and WASP 2.0

In this section we compare WASP 2.0 to WASP 1.0. First of all we observe that WASP 1.0 does not implement any program transformation phase, whereas WASP 2.0 applies both Clark's completion and program simplification in the style of (Eén and Biere 2005). The addition of this preprocessing step brings advantages in both terms of simplifying the implementation of the propagation procedure and in terms performance. The Clark's completion introduces a number of clauses that represent support propagation, which is implemented natively in WASP 1.0 instead. The subsequent program simplification step optimizes the program by eliminating redundant atoms (also introduced by the completion) and shrinking definitions. This results in a program that is usually easier to evaluate. Concerning the well-founded operator both WASP 2.0 and WASP 1.0 compute unfounded sets according to the *source pointers* (Simons, Niemelä, and Soininen 2002) technique. WASP 1.0, which implements a native inference rule, immediately infers unfounded atoms as false, and updates a special implementation of the implication graph. In contrast, WASP 2.0 learns a clause representing the inference (also called loop formula) and propagates it with unit propagation. This choice combined with Clark's completion allows to simplify conflict analysis, learning and backjumping. Indeed, WASP 1.0 implements specialized variants of these procedures that require the usage of complex data structures that are difficult to optimize. Since in WASP 2.0 literals are always inferred by the UnitPropagation procedure, we could adopt an implementation of these strategies optimized as in modern SAT solvers. Finally both WASP 2.0 and WASP 1.0 implement conflict-driven branching heuristics. WASP 2.0 uses a branching heuristic inspired to the one of MiniSAT, while WASP 1.0 uses an extension of the BerkMin (Goldberg and Novikov 2002) heuristics extended by adding a look-ahead technique and an additional ASP-specific criterion.

Experiment

In this section we report the results of an experiment assessing the performance of WASP 2.0. In particular, we compare WASP 2.0 with WASP 1.0 and clasp. All the solvers used gringo 3.0.5 (Gebser et al. 2011) as grounder. clasp and WASP 1.0 has been executed with the same heuristic setting used in (Alviano et al. 2013). Concerning clasp we used the version 3.0.1. The experiment was run on a Mac Pro equipped with two 3 GHz Intel Xeon X5365 (quad core) processors, with 4 MB of L2 cache and 16 GB of RAM, running Debian Linux 7.3 (kernel ver. 3.2.0-4-amd64). Binaries were generated with the GNU C++ compiler 4.7.3-4 shipped by Debian. Time limit was set to 600 seconds. Performance was measured using the tools pyrunlim and pyrunner (*https://github.com/alviano/python*).

Tested instances are among those in the System Track of the 3rd ASP Competition (Calimeri, Ianni, and Ricca 2014), in particular all instances in the NP category. This category includes planning domains, temporal and spatial scheduling problems, combinatorial puzzles, graph problems, and a number of real-world domains in which ASP has been applied. (See (Calimeri, Ianni, and Ricca 2014) for an exhaustive description of the benchmarks.)

Table 1 summarizes the number of solved instances and the average running times in seconds for each solver. In particular, the first two columns report the total number of instances (#) and the number of instances that are solved by all solvers ($\#_{all}$), respectively; the remaining columns report the number of solved instances within the time-out (sol.), and the running times averaged both over solved instances (t) and over instances solved by all variants (t_{all}).

Table 1: Average running time and number of solved instances

			clasp			WASP 1.0			WASP 2.0		
Problem	**#**	**#$_{all}$**	**sol.**	**t**	**t$_{all}$**	**sol.**	**t**	**t$_{all}$**	**sol.**	**t**	**t$_{all}$**
DisjunctiveScheduling	10	5	5	16.8	16.8	5	29.0	29.0	5	188.4	188.4
GraphColouring	10	3	4	88.0	20.6	3	50.5	50.5	3	3.3	3.3
HanoiTower	10	2	7	126.0	49.8	2	214.0	214.0	7	52.5	18.3
KnightTour	10	6	10	14.3	0.3	6	93.5	93.5	10	16.0	0.6
Labyrinth	10	8	9	74.4	74.7	8	118.7	118.7	10	85.8	84.7
MazeGeneration	10	10	10	0.3	0.3	10	19.9	19.9	10	2.7	2.7
MultiContextSystemQuerying	10	10	10	5.1	5.1	10	122.4	122.4	10	9.4	9.4
Numberlink	10	6	8	21.1	0.6	6	24.3	24.3	7	8.7	5.5
PackingProblem	10	0	0	-	-	0	-	-	0	-	-
SokobanDecision	10	5	10	101.5	2.8	5	212.8	212.8	7	97.8	14.4
Solitaire	10	2	2	124.9	124.9	3	183.1	198.0	4	8.7	6.0
WeightAssignmentTree	10	1	5	119.2	22.4	1	297.3	297.3	3	282.3	97.9
Total	**120**	**58**	**80**	**62.9**	**20.5**	**59**	**124.1**	**95.6**	**76**	**68.7**	**34.6**

We observe that WASP 2.0 outperforms WASP 1.0. In fact, WASP 2.0 solved 17 instances more than WASP 1.0, and also the improvement on the average execution time is sensible, with a percentage gain of around 64% on instances solved by all systems. On the other hand, clasp is faster than WASP 2.0, with a percentage gain of around 41 % on the same instances. Moreover, clasp solved 4 instances more than WASP 2.0.

Analyzing the results in more detail, there are some specific benchmarks where WASP 2.0 and clasp exhibit significantly performances. Two of these problems are SokobanDecision and WeightAssignmentTree, where clasp solved 3 and 2 instances more than WASP 2.0, respectively, while WASP 2.0 solved 2 instances more than clasp in Solitaire. We also note that the performance of WASP deteriored in DisjunctiveScheduling. This is due to the initial steps of the computation, and in particular to the simplification procedure, which in this case removes 80% of clauses and 99% of atoms. However, there are cases in which simplifications play a crucial role to improve performance of the answer set search procedure. For example, in HanoiTower, where WASP 2.0 performs better than other systems, more than half of the variables are removed in a few seconds.

Related Work

WASP 1.0 is inspired by several techniques used in SAT solving that were first introduced for Constraint Satisfaction and QBF solving.

Some of these techniques were already adapted in non-disjunctive ASP solvers like Smodels$_{cc}$ (Ward and Schlipf 2004), clasp (Gebser et al. 2007), Smodels (Simons, Niemelä, and Soininen 2002), Cmodels3 (Lierler and Maratea 2004), and DLV (Ricca, Faber, and Leone 2006). More in detail, WASP 2.0 differs from Cmodels3 (Lierler and Maratea 2004) that are based on a rewriting into a propositional formula and an external SAT solver. WASP 2.0 differs from DLV (Alviano et al. 2011) and the Smodels variants, which features a native implementation of all inference rules. Our new solver is more similar to clasp, but there are differences concerning the restart policy, constraint deletion and branching heuristics. WASP 2.0 adopts as default a policy based on the sequence of thresholds introduced in (Luby, Sinclair, and Zuckerman 1993), whereas clasp employs by default a different policy based on geometric series. Concerning deletion of learned constraints, WASP 2.0 adopts a criterion inspired by MiniSAT, while clasp implements a technique introduced in Glucose (Audemard and Simon 2009). Moreover, clasp adopts a branching heuristic based on BerkMin (Goldberg and Novikov 2002) with a variant of the MOMS criterion which estimates the effect of the candidate literals in short clauses.

Conclusion

In this paper we reported on the recent improvement of the ASP solver WASP 1.0. We described the main improvements on the evaluation procedure focusing on the improvements to the core evaluation algorithms working on normal programs. The new solver was compared with both its predecessor and the latest version of clasp on on benchmarks belonging to the NP class, where WASP 1.0 was not competitive. The result is very encouraging, since WASP 2.0 improves substantially w.r.t. WASP 1.0 and is often competitive with clasp.

Future work concerns the reengineering of disjunctive rules, aggregates, and weak constraints, as well as the introduction of a native implementation of choice rules.

References

[Alviano et al. 2011] Alviano, M.; Faber, W.; Leone, N.; Perri, S.; Pfeifer, G.; and Terracina, G. 2011. The disjunctive datalog system DLV. In Gottlob, G., ed., *Datalog 2.0*, volume 6702. Springer Berlin/Heidelberg. 282–301.

[Alviano et al. 2013] Alviano, M.; Dodaro, C.; Faber, W.; Leone, N.; and Ricca, F. 2013. Wasp: A native asp solver based on constraint learning. In Cabalar, P., and Son, T. C., eds., *LPNMR*, volume 8148 of *LNCS*, 54–66. Springer.

[Audemard and Simon 2009] Audemard, G., and Simon, L. 2009. Predicting learnt clauses quality in modern sat solvers. In Boutilier, C., ed., *IJCAI*, 399–404.

[Calimeri, Ianni, and Ricca 2014] Calimeri, F.; Ianni, G.; and Ricca, F. 2014. The third open answer set programming competition. *Theory and Practice of Logic Programming* 14(1):117–135.

[Davis, Logemann, and Loveland 1962] Davis, M.; Logemann, G.; and Loveland, D. 1962. A Machine Program for Theorem Proving. *Communications of the ACM* 5:394–397.

[Eén and Biere 2005] Eén, N., and Biere, A. 2005. Effective preprocessing in sat through variable and clause elimination. In *SAT*, volume 3569 of *LNCS*, 61–75. Springer.

[Eén and Sörensson 2003] Eén, N., and Sörensson, N. 2003. An extensible sat-solver. In Giunchiglia, E., and Tacchella, A., eds., *SAT*, volume 2919 of *LNCS*, 502–518. Springer.

[Gaschnig 1979] Gaschnig, J. 1979. *Performance measurement and analysis of certain search algorithms*. Ph.D. Dissertation, Carnegie Mellon University, Pittsburgh, PA, USA. Technical Report CMU-CS-79-124.

[Gebser et al. 2007] Gebser, M.; Kaufmann, B.; Neumann, A.; and Schaub, T. 2007. Conflict-driven answer set solving. In *IJCAI*, 386–392. Morgan Kaufmann Publishers.

[Gebser et al. 2011] Gebser, M.; Kaminski, R.; König, A.; and Schaub, T. 2011. Advances in *gringo* series 3. In Delgrande, J. P., and Faber, W., eds., *LPNMR*, volume 6645 of *LNCS*, 345–351. Springer.

[Gelfond and Lifschitz 1991] Gelfond, M., and Lifschitz, V. 1991. Classical Negation in Logic Programs and Disjunctive Databases. *New Generation Computing* 9:365–385.

[Goldberg and Novikov 2002] Goldberg, E., and Novikov, Y. 2002. BerkMin: A Fast and Robust Sat-Solver. In *Design, Automation and Test in Europe Conference and Exposition (DATE 2002)*, 142–149. Paris, France: IEEE Computer Society.

[Gomes, Selman, and Kautz 1998] Gomes, C. P.; Selman, B.; and Kautz, H. A. 1998. Boosting Combinatorial Search Through Randomization. In *Proceedings of AAAI/IAAI 1998*, 431–437. AAAI Press.

[Lierler and Maratea 2004] Lierler, Y., and Maratea, M. 2004. Cmodels-2: SAT-based Answer Set Solver Enhanced to Non-tight Programs. In Lifschitz, V., and Niemelä, I., eds., *Proceedings of LPNMR*, volume 2923 of *LNAI*, 346–350. Springer.

[Luby, Sinclair, and Zuckerman 1993] Luby, M.; Sinclair, A.; and Zuckerman, D. 1993. Optimal speedup of las vegas algorithms. *Inf. Process. Lett.* 47:173–180.

[Moskewicz et al. 2001] Moskewicz, M. W.; Madigan, C. F.; Zhao, Y.; Zhang, L.; and Malik, S. 2001. Chaff: Engineering an Efficient SAT Solver. In *Proceedings of the 38th DAC*, 530–535. Las Vegas, NV, USA: ACM.

[Ricca, Faber, and Leone 2006] Ricca, F.; Faber, W.; and Leone, N. 2006. A Backjumping Technique for Disjunctive Logic Programming. 19(2):155–172.

[Simons, Niemelä, and Soininen 2002] Simons, P.; Niemelä, I.; and Soininen, T. 2002. Extending and Implementing the Stable Model Semantics. *Artificial Intelligence* 138:181–234.

[Ward and Schlipf 2004] Ward, J., and Schlipf, J. S. 2004. Answer Set Programming with Clause Learning. In Lifschitz, V., and Niemelä, I., eds., *Proceedings of the 7th International Conference on Logic Programming and Non-Monotonic Reasoning (LPNMR-7)*, volume 2923 of *LNAI*, 302–313. Springer.

[Zhang et al. 2001] Zhang, L.; Madigan, C. F.; Moskewicz, M. W.; and Malik, S. 2001. Efficient Conflict Driven Learning in Boolean Satisfiability Solver. In *Proceedings of the ICCAD*, 279–285.